UNITED STATES PATENT APPLICATION

OF

KEN KUWABARA

STEVEN LIN

AND

MICHAEL LANGDON

FOR

FILTER-BASED FORWARDING IN A NETWORK

FILTER-BASED FORWARDING IN A NETWORK

BACKGROUND OF THE INVENTION

[0001] In a conventional computer network (e.g., the Internet), computers communicate over a network infrastructure made up of interconnected nodes, such as routers and/or switches, connected by communication links such as optical fiber, copper cable, and wireless links. Typically, the topology of the network infrastructure is configured in such a way that the infrastructure contains more than one path through which data may be carried from one computer to another. The topology, for example, may include a partial mesh configuration, where a node is connected to multiple other nodes. A router in such a network, therefore, may contain a plurality of interface ports for connection to multiple neighboring routers.

[0002] Such a router typically receives data in discrete units (herein referred to as "packets," which may include frames, cells, packets, or any other fixed- or variable-sized unit of data) at one or more of its ingress interface ports. The router examines destination address information embedded in the packets and determines the appropriate egress interface ports for outputting the respective packets, typically by performing a table lookup. To construct and update tables, routers may use dynamic routing protocols to systematically exchange information with other devices in the network to obtain a view of the network topology (this information being maintained in a routing database, such as one or more routing tables). Based on this information, the router constructs and updates a forwarding table, which associates ranges of destination addresses to respective egress interface ports.

[0003] In some cases, however, such use of forwarding tables may be inadequate. By relying on destination addresses to determine the appropriate egress interface port for packets,

traditional routers do not distinguish packets according to other criteria. It may be desirable to use other criteria to, for example, facilitate traffic engineering of certain types of packets (i.e., select egress interface ports based on packet type as well as destination address).

[0004] In addition, in certain circumstances, the traditional use of forwarding tables may be inadequate to implement virtual private networks (VPNs). In cases where a single router forwards traffic for two separate VPNs, the router needs to ensure that traffic from one VPN is not sent to the other VPN. One proposed solution is to bind one or more ingress interface ports and one or more egress interface ports to each VPN. In this way, the bound ingress and egress interface ports only carry traffic for one VPN, allowing the router to readily maintain separation of traffic for each VPN.

[0005] In some cases, however, the network configuration may be such that traffic from two VPNs is intermingled and received at a single ingress interface port of a router. This may happen, for example, where the traffic from the two VPNs is carried over an open access network in which traffic separation is not maintained before arriving at the router. One potential solution is to inject tags into each packet to uniquely identify the VPN from which the packets came. This may be undesirable because additional components or enhancements would be required at each source computer in the VPN to generate tags and at the router to identify the tags and separate the different VPN traffic. Another solution is to use policy-based routing, which involves statically configuring the forwarding table to forward packets according to criteria other than destination address. This may also be undesirable because static policies are typically configured manually and are not updated dynamically as the state of the network changes. Policy-based routing may also require the use of additional components or enhancements.

[0006] Thus, there is a need for an invention that more adequately addresses problems occurring in the network.

SUMMARY OF THE INVENTION

[0007] According to one embodiment of the invention, a router receives a packet at an ingress interface. The router classifies the received packet based on at least a first field value contained in the header of the packet. According to the classification of the received packet, the router associates one of the plurality of forwarding tables to the packet. The router then performs a lookup operation in the associated forwarding table according to at least a second field value contained in the header of the packet. Based on the lookup operation, the router determines an egress interface and transmits the received packet from the determined egress interface.

[0008] It is to be understood that both the foregoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the invention, as claimed.

BRIEF DESCRIPTION OF THE DRAWINGS

[0009] The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate several embodiments of the invention and together with the description, serve to explain the principles of the invention.

[0010] Figure 1A is a block diagram of a router incorporating principles of the invention according to a first embodiment;

[0011] Figure 1B is a block diagram of a router incorporating principles of the invention according to a second embodiment;

[0012] Figure 2 is a block diagram of a route lookup module incorporating principles of the invention;

[0013] Figure 3 is a flow chart representing the initialization process for a router according to the invention;

[0014] Figure 4 is a flow diagram representing a process for forwarding packets in a router according to the invention;

[0015] Figure 5 shows a first example of a network topology using a router according to the invention;

[0016] Figure 6 shows a second example of a network topology using a router according to the invention;

[0017] Figure 7 shows a third example of a network topology using a router according to the invention; and

[0018] Figure 8 shows a fourth example of a network topology using a router according to the invention.

DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0019] Reference will now be made in detail to embodiments of the invention, examples of which are illustrated in the accompanying drawings. Wherever possible, the same reference numbers will be used throughout the drawings to refer to the same or like parts.

[0020] According to the invention, a router may perform packet forwarding based on not only the destination address of the packets, but also based upon other information contained in the packets. Such other information may include source address, protocol field, packet classification, and packet type. In one embodiment, the router of the present invention uses a filter, such as a firewall filter, to classify packets based on packet header field values and selects respective forwarding tables for lookup based on the classification.

[0021] Figure 1A shows router 100. Generally, router 100 receives packets on the input lines and determines the output lines from which the packets are to be output. As shown in Figure 1A, router 100 includes a plurality of interface (I/F) modules 110, switch fabric(s) 120, route lookup module(s) 130, and routing engine 140. Each I/F module 110 connects to one or more respective input lines and/or output lines.

[0022] Each I/F module 110 contains one or more ingress interface ports (not shown) for receiving packets from respective input lines and/or one or more egress interface ports (not shown) for transmitting packets to respective output lines. I/F module 110 may perform processing on the headers of respective packets (e.g., layer 2/3 header processing) so the packets are in the appropriate format for processing through switch fabric 120 (for packets received at the ingress interface ports) and for transmission over the output lines (for packets to be sent out the egress interface ports). In one embodiment, I/F modules 110 are implemented as circuit boards that are insertable and removable from router 100. One or more I/F modules 110 may be inserted at a time.

[0023] Switch fabric 120 contains logic for receiving packets from I/F modules 110 and for transmitting packets out to I/F modules 110. In one embodiment, switch fabric 120 may contain buffer memory for storing received packets. In another embodiment, switch fabric may be a cross-connect connecting sets of I/F modules 110 with buffer memory contained in respective I/F modules 110. Switch fabric 120 preferably provides packet header information (e.g., destination address information) for received packets to route lookup module 130 and receives identifiers specifying which egress interface ports in the I/F modules 110 to forward the packets to.

[0024] Each of route lookup module(s) 130 examines characteristics of packets flowing through switch fabric 120 to determine the manner in which they are to be forwarded (e.g., determining the egress interface ports on which respective packets are to be output). In one embodiment, route lookup module 130 examines packet header information for respective packets, which may include performing a filtering operation and table lookups using one or more forwarding tables. One embodiment of route lookup module 130 is described in connection with Figure 2.

[0025] Routing engine 140 maintains a database of routing information which may be received in route messages from other routers using routing protocols. Routing engine 140 preferably generates forwarding tables which are transmitted to route lookup module 130. Routing engine 140 systematically receives updated information reflecting changes in the network, updates its database (which may be stored as routing tables), generates one or more forwarding tables from the database, and transmits the forwarding tables to route lookup module 130. Routing engine 140 preferably comprises a processor executing routing protocols for computing routes and/or network topology.

[0026] Figure 1B is a block diagram of router 102. Router 102 is similar to router 100, except that each I/F module 110 has a corresponding route lookup module 130 connected to it. In addition, routing engine 140 connects to each of the route lookup modules 130. Router 102 operates in a manner similar to router 100, except that a lookup operation is performed by the route lookup module 130 corresponding to the I/F module 110 at which packets are received. I/F modules 110 provide packet header information to the corresponding route lookup modules 130 and receive egress interface port information for each packet. I/F modules 110 transmit the packets and received egress interface port information to switch fabric 120, which forwards the

packets to the appropriate I/F modules 110. To carry out this operation, each I/F module 110 in router 102 may contain buffer memory for storing packets.

[0027] Figures 1A and 1B show different embodiments of routers that may be used with the invention. In alternative embodiments, the invention may be used with other router or switch architectures or in any device that performs packet forwarding.

[0028] Figure 2 shows a block diagram of route lookup module 130. As shown in Figure 2, route lookup module 130 includes filter 200, lookup processor 210, and a plurality of forwarding tables 220 (shown as 220(A) and 220(B)) stored in memory. While two forwarding tables 220 are shown in Figure 2, any number of forwarding tables may be used in alternative embodiments. Lookup processor 210 is connected to filter 200 and forwarding tables 220.

[0029] Filter 200 is preferably a firewall filter which may be programmed to classify or identify packets based on selected criteria and perform certain actions based on those classifications. Such criteria, for example, may involve the content of received packets, such as IP destination address, IP source address, IP protocol field, the ingress and/or egress router interfaces, and the state of the router. In a preferred embodiment, filter 200 classifies packets based on source address information and selects one of forwarding tables 220 based on that classification. In other embodiments, other actions may be carried out in addition to the selection of one of forwarding tables 220. Filter 200 may be implemented in hardware as circuit logic for carrying out the respective operation or as one or more processors programmed to carry out the operation.

[0030] Lookup processor 210 performs table lookups. For each packet, lookup processor 210 preferably receives packet header information, including the destination addresses of the packets, and a table identifier specifying a particular forwarding table 220 to be used and outputs

an egress interface port identifier for that packet. In addition to an egress interface port identifier, lookup processor 210 may, in appropriate cases, output the address of the neighboring device to which the packet is to be sent. This may be appropriate, for example, when the egress interface port is an Ethernet port. Lookup processor 210 may be implemented in hardware as circuit logic for carrying out the respective operation or as one or more processors programmed to carry out the operation.

[0031] Each of forwarding tables 220 contains entries associating ranges of destination addresses to corresponding egress interface ports in the router. In one embodiment, each forwarding table 220 corresponds to a separate virtual private network and at least some of the entries in each forwarding table 220 correspond to label switched paths that have been established to other nodes in the respective virtual private network.

[0032] While forwarding tables 220 may contain entries associating address ranges with respective egress interface port identifiers, alternative implementations of forwarding tables 220 may contain one or more levels of indirection. For example, forwarding tables 220 may associate address ranges with nexthop identifiers. The nexthop identifiers are associated with egress interface ports (or other actions) in a nexthop resolution table.

[0033] Figure 3 shows a flow diagram of a process for initializing a router in accordance with the invention. A forwarding table is generated for each classification that will be programmed into the filter (step 310). The filter is then programmed to define the criteria used for packet classification (step 320). Such classifications may, for example, include ranges of source addresses, other packet header criteria, and a default if no criteria is met. The filter is programmed to select a forwarding table corresponding to each respective classification (step 330).

[0034] Figure 4 shows a flow diagram representing a process for the operation of a router (such as router 100 or 102) in accordance with the present invention. A packet received at the router is classified based on packet header information (step 410). This step may be performed by a programmed filter, such as a firewall filter that uses packet header information criteria to classify the packet. In a preferred embodiment, the filter classifies packets based on source address. The packet may be classified in the default classification if no criteria is met for that packet. A forwarding table is then selected based upon the classification (step 420). In a preferred embodiment, the filter selects a forwarding table for a packet based on the classification of that packet.

[0035] In alternative embodiments, additional or alternative actions and/or packet processing may be performed based upon the classification (step 430). Such other actions may include, for example, sampling, policing, logging, and setting alerts. Packet processing that may be performed may include label encapsulation and/or decapsulation. The output port on which the packet will be output is determined using the selected forwarding table (step 440). This may be done by using the packet header information, such as the destination addresses, to look up the corresponding egress interface output port identifier (and/or neighboring device address) in the selected forwarding table.

[0036] Figure 5 shows an example of a network topology using a router of the present invention. Two Internet Service Providers (ISPs) are represented: ISP A and ISP B. Each ISP customer desires to access its respective ISP's network and must do so through an open-access network. Router F serves as a gateway between the open-access network and each ISP network.

[0037] When data from the ISP A customer and the ISP B customer are sent to the open access network, router F operates to separate those packets and transmit them to the networks of

the respective ISPs. Router F, for example, may be programmed such that its firewall filter detects packets coming from the ISP A customer by identifying the source addresses of packets from the ISP A customer and detects packets from the ISP B customer by identifying the source addresses of packets from the ISP B customer. When the firewall filter in router F determines that a packet has come from either the ISP A customer or the ISP B customer, it can then select the appropriate forwarding table which contains entries to egress interface ports connected to the respective ISP network.

[0038] Figure 6 shows a network topology similar to the one shown in Figure 5, further including a VPN backbone connecting router F to routers A2 and B2. Established within the VPN backbone are LSP 1 and LSP 2 over which packets destined to the networks of ISP A and ISP B are transmitted, respectively. As in the network shown in Figure 5, router F classifies packets using its firewall filter to identify which packets have come from the ISP A customer and which packets have come from the ISP B customer. Router F uses the appropriate forwarding tables based on its classifications of the packets. Those forwarding tables may contain entries corresponding to respective LSPs 1 and 2.

[0039] Figure 7 shows an alternative network topology to the one shown in Figure 6. As shown in Figure 7, two routers, routers F1 and F2, are used in place of router F. Router F1 performs filtering and separates packets from the ISP A customer and the ISP B customer. Packets from each of these respective ISP customers are transmitted over separate egress interface ports of router F1 and accordingly sent over separate links to router F2. In router F2, sets of ingress and egress interface ports are bound to the networks of respective ISPs. That is, traffic received at one ingress interface port is forwarded out of the egress interface port to which it is bound. Each set of interface ports corresponds to an ISP.

[0040] Figure 8 shows another exemplary network topology that utilizes a router in accordance with the present invention. The network topology shown in Figure 8 illustrates how the router of the present invention may be used for traffic engineering. In a typical case, packets coming into router F would be routed to router 3 based on an algorithm that determines a least cost path. Here, packets coming into router F would be routed to router 3 via router 1 (1+1=2) because the combined metric through router 1 is lower than the combined metric through router 2 (1+2=3). To facilitate traffic engineering, the firewall filter in router F may be programmed to classify certain packets received by router F, such as voice traffic or traffic from certain source addresses, and select a routing table that would forward such packets to router 3 via router 2. In this way, certain kinds of traffic can be routed over a different path then what would otherwise be computed using least cost path computations.

[0041] Other embodiments of the invention will be apparent to those skilled in the art from consideration of the specification and practice of the invention disclosed herein. For example, while the invention has been described in connection with a router, the invention may also be used in a switch or other networking device in which actions are performed on packets. Further, while the invention has been described in connection with tables (routing tables and forwarding tables), the invention may also use route or forwarding information stored in other data structures/forms or in databases. It is intended that the specification and examples be considered as exemplary only, with a true scope and spirit of the invention being indicated by the following claims.